

Statistical physics approach to graphical games: local and global interactions

A. Ramezanpour* and J. Realpe-Gomez†

Politecnico di Torino, C.so Duca degli Abruzzi 24, I-10129 Torino, Italy

R. Zecchina‡

Politecnico di Torino, C.so Duca degli Abruzzi 24, I-10129 Torino, Italy and

Collegio Carlo Alberto, Via Real Collegio 30, 10024 Moncalieri, Italy

(Dated: January 19, 2013)

Abstract

In a graphical game agents play with their neighbors on a graph to achieve an appropriate state of equilibrium. Here relevant problems are characterizing the equilibrium set and discovering efficient algorithms to find such an equilibrium (solution). We consider a representation of games that extends over graphical games to deal conveniently with both local and global interactions and use the cavity method of statistical physics to study the geometrical structure of the equilibria space. The method also provides a distributive and local algorithm to find an equilibrium. For simplicity we consider only pure Nash equilibria but the methods can as well be extended to deal with (approximated) mixed Nash equilibria.

* abolfazl.ramezanpour@polito.it

† john.realpe@polito.it

‡ riccardo.zecchina@polito.it

I. INTRODUCTION

In the last decade we observed a rapid merging of research interests in social sciences, economics and computer science driven in part by the common need of analyzing strategic interactions in multi-agent systems. There is a growing body of empirical work evidencing the influence of social interactions on economic outcomes which has encouraged economists to take explicit account of the *direct* (non-market) social influences on individual decision making [1–3].

A general framework to study strategic interactions in multi-agent systems is game theory[4]. It serves to analyze situations where self-interested agents, possibly with conflicting interests, struggle to get the best conditioned on (and conditioning) the behavior of others; the scope of the theory is to predict the strategic behavior that comes out in such situations, something that is usually done in the form of equilibrium concepts. Chief among them is the concept of Nash equilibrium (NE) that is a strategy profile in which no agent has incentives to deviate *unilaterally*. Nash equilibria can be pure or mixed, depending on whether the agents are required to play deterministically or are allowed to randomize among their available strategies. While conceptually very compelling, the concept of pure NE has the drawback that it does not always exist, in contrast with the universality of the concept of mixed NE whose existence is guaranteed for any finite game [5, 6].

An important concern for computer scientists was to understand if Nash equilibria are actually efficiently computable [7]. Much work has been devoted at analyzing the computational complexity of finding a Nash equilibrium [8–11]. This has been proved to be complete in a class called PPAD [12] which contains problems believed to be hard problems[13, 14]. Indeed finding a Nash equilibrium typically becomes NP-hard as soon as we require it to satisfy certain natural properties (e.g. optimizing social welfare) [15, 16] or when we restrict it to pure strategies [17, 18].

Statistical physics has also contributed important insights and techniques to these fields [19]. For instance, in computer science it has provided powerful heuristic algorithms and a better understanding of the onset of computational complexity. From a physical point of view, a game is regarded as a system of interacting agents where an appropriate energy function maps the Nash equilibria to the ground states of the system. Our understanding of these systems has considerably improved in recent years mainly due to the concepts and

tools developed in the study of complex systems displaying glassy behaviors [20–23].

From an economics modeling perspective, an interesting aspect of graphical games is to study the interplay between local and global interactions [24–26]. Motivated by this, we will study a global graphical game where agents, besides the local payoffs, receive some global payoffs depending on an aggregate quantity, here the average strategy of the game or global magnetization. We design a message passing algorithm with a polynomial time complexity which is exact when the graph of local interactions is a tree and show how to obtain an approximated algorithm to deal with the global interaction in a more efficient way. This algorithm resembles the two step strategy used by Horst and Scheinkman [24] to prove the existence of equilibria in a class of multi-agent systems with local (pairwise) and global interactions. Namely, we consider the magnetization as a fixed parameter, rendering the game an effective local graphical game, and require consistency in that we force the average magnetization to be equal to the fixed parameter. We test all these algorithms in some ensembles of random graphical games with global interactions, focusing mainly on an extension of the best shot game or maximal independent set problem [27, 28]. Here we present upper bounds for the entropy and for the probability of having solutions which along with the results obtained by the cavity method of statistical physics help us to characterize the solution space of the problem.

In this paper we take the constraint satisfaction approach and for the sake of simplicity we focus on pure Nash equilibria, but the methods can properly be generalized to deal with mixed equilibria. We show numerically how the information contained in the Belief Propagation (BP) messages [29] can be exploited to turn the algorithm into a one phase, fully distributed (and typically efficient) solver which directly converges to a *single* Nash equilibrium. This class of message passing solvers are called reinforced belief propagation (rBP)[30, 31]. Here we will consider the ensemble of independent payoffs that could be fully random or with some hidden solutions; in both cases we find typical Nash equilibria, if they exist, in a replica symmetric phase and easy to discover with the rBP algorithm. However, finding an optimal Nash equilibrium maximizing the total payoff would be still difficult as we enter into a replica symmetry broken phase with a more complex organization of the optimal solutions.

The paper is organized as follows. In section II we give some definitions that will be used in the paper. Section III includes our results on graphical games with local random payoffs

which could have hidden solutions. In section IV we study graphical games with a global interaction that depends on the average strategy of the agents. The conclusion is given in section V. In appendix A we present more details on some rigorous statements mentioned in the text.

II. RELATED WORKS AND DEFINITIONS

Traditionally, a game is defined by assigning a number (i.e. a payoff) to each player for each possible configuration of the other players. The number of parameters involved in this representation grows exponentially with the number of agents in the game. Therefore, it is crucial to exploit additional structures that may be present in certain situations to find an efficient representation of the problem. Graphical games model the very common situation where each agent interacts only with a small subset of the whole population (see [32–35] for related works). The interaction structure is encoded in a graph where each vertex corresponds to an agent and a link between two agents indicates that the two player’s payoffs depend on each other’s choice of strategy.

In Ref. [34] the authors provide a dynamic algorithm to find Nash equilibria on trees, which was later extended to deal with loopy graphs [36]. They also drew analogies with the belief propagation algorithm [29] which is equivalent to the replica symmetric approximation in the cavity method. In fact the first phase of their algorithm can be understood as an instance of Warning Propagation (WP) where the probabilistic information contained in a belief is projected onto a boolean variable. Such a simplification allows to keep things in the realm of integer arithmetic and, in the case of graphical games, to prove the algorithm convergence. Connections to constraint satisfaction problems and Markov random fields can be found in [37, 38].

Consider N players indexed by $i = 1, \dots, N$ playing a game with given payoffs $M_i(\sigma_i|\sigma_{\partial i})$. A player payoff or utility depends on her strategy $\sigma_i \in \Lambda_i$ and the set of strategies played by her neighbors in a dependency graph \mathcal{G} , i.e. $\sigma_{\partial i}$. We use ∂i to denote the neighbor set of player i . In the following we consider binary strategies, that is $\Lambda_i = \{-1, +1\}$, and positive payoffs $M_i(\sigma_i|\sigma_{\partial i}) \in [0, 1]$.

A strategy profile or configuration $\underline{\sigma}^*$ is a pure Nash equilibrium if for each player σ_i^* is the best response to the neighbors actions, i.e.

$$M_i(\sigma_i^*|\sigma_{\partial i}^*) = \max_{\sigma_i} M_i(\sigma_i|\sigma_{\partial i}^*). \quad (1)$$

Not any game possesses a pure equilibrium, but Nash theorem ensures that any finite game admits a mixed equilibrium [6]. Consider stochastic players where player i chooses strategy $\sigma_i = +1$ with probability $x_i \in [0, 1]$ and the other strategy with the complement probability. A mixed profile \underline{x}^* is called a mixed Nash equilibrium if

$$\langle M_i(\sigma_i|\sigma_{\partial i}) \rangle_{\sigma_i, \sigma_{\partial i}} = \max_{\sigma_i} \langle M_i(\sigma_i|\sigma_{\partial i}) \rangle_{\sigma_{\partial i}}, \quad (2)$$

where the averages are taken over strategies with respect to \underline{x}^* .

In this paper we shall mainly consider pure Nash equilibria. In cases where such an equilibrium does not exist, we ask for approximate solutions that satisfy Nash conditions within some tolerated errors [7]. A configuration $\underline{\sigma}^*$ is called an ϵ -Nash equilibrium if

$$M_i(\sigma_i^*|\sigma_{\partial i}^*) \geq M_i(-\sigma_i^*|\sigma_{\partial i}^*) - \epsilon. \quad (3)$$

where $\epsilon \geq 0$.

We shall represent the above problem as a constraint satisfaction problem where a strategy profile is a Nash equilibrium if and only if it satisfies all the constraints. To this end, we define a constraint $I_i(\sigma_i|\sigma_{\partial i})$ for each player i to check if strategy σ_i is a best response or not. That is $I_i(\sigma_i|\sigma_{\partial i}) = 1$ if σ_i is the best response to the neighborhood configuration $\sigma_{\partial i}$, otherwise it is zero. Statistical properties of the solution space can be obtained by studying the following partition function

$$Z = \sum_{\underline{\sigma}} e^{\beta \sum_i M_i(\sigma_i|\sigma_{\partial i})} \prod_i I_i(\sigma_i|\sigma_{\partial i}), \quad (4)$$

where β is a parameter to optimize over the solution space. Two special limits $\beta = 0$ and $\beta \rightarrow \infty$ give the typical and optimal welfare solutions, respectively.

The reader can find more about physical approaches to games in [39, 40] and references therein.

III. LOCAL GRAPHICAL GAMES

We first study local graphical games where each player payoff depends only on its local neighborhood on the graph. In particular we consider random regular graphs ($|\partial i| = K$ for all players). The aim is to characterize the equilibria space, for example the number of equilibria and their geometrical organization. To do this, we resort to the cavity method of statistical physics in the replica symmetric (RS) and 1-step replica symmetry breaking (1RSB) approximations [19].

A. RS equations

In the replica symmetric approximation we assume that all solutions belong to a single cluster in the configuration space; there is a path connecting any two solutions such that neighboring solutions along the path are different only on a sub-linear number of players. As a result, correlations are short ranged and it is usually computationally easy to find a solution (a Nash equilibrium) to the problem.

Assume that our dependency graph \mathcal{G} is a tree and consider cavity graph $\mathcal{G}_{i \rightarrow j}$ including all the nodes and edges connected to j through i . Define $\pi_{i \rightarrow j}(\sigma_i; \sigma_j)$ as the probability of having strategies σ_i and σ_j in a solution of $\mathcal{G}_{i \rightarrow j}$ when constraint I_j is ignored. Then one can easily write an equation for $\pi_{i \rightarrow j}(\sigma_i; \sigma_j)$ relating it to other cavity probabilities [19]:

$$\pi_{i \rightarrow j}(\sigma_i; \sigma_j) \propto \sum_{\{\sigma_k | k \in \partial i \setminus j\}} e^{\beta M_i I_i} \prod_{k \in \partial i \setminus j} \pi_{k \rightarrow i}(\sigma_k; \sigma_i). \quad (5)$$

These equations are called belief propagation equations and can be solved iteratively starting from random initial cavity probabilities or messages. Having the cavity probabilities we can obtain the free energy by the Bethe expression:

$$F = \sum_i \Delta F_i - \sum_{(ij) \in \mathcal{G}} \Delta F_{ij}, \quad (6)$$

where ΔF_i and ΔF_{ij} are the free energy shifts by adding node i and link (ij) , respectively. In a tree graph one obtains:

$$\begin{aligned}
e^{-\beta\Delta F_i} &= \sum_{\sigma_i} \sum_{\sigma_{\partial i}} e^{\beta M_i I_i} \prod_{j \in \partial i} \pi_{j \rightarrow i}(\sigma_j; \sigma_i), \\
e^{-\beta\Delta F_{ij}} &= \sum_{\sigma_i, \sigma_j} \pi_{i \rightarrow j}(\sigma_i; \sigma_j) \pi_{j \rightarrow i}(\sigma_j; \sigma_i).
\end{aligned} \tag{7}$$

We expect the free energy to be asymptotically correct in locally tree graphs as long as correlations are short range, i.e. in a replica symmetric phase.

The BP equations can also be used as an algorithm to find a Nash equilibrium. Reinforcement is a way of doing this by progressively biasing the players to take the strategy that is suggested by the BP marginals [31]. The reinforced BP equations are

$$\begin{aligned}
\pi_{i \rightarrow j}(\sigma_i; \sigma_j) &\propto [\pi_i(\sigma_i)]^r \sum_{\{\sigma_k | k \in \partial i \setminus j\}} e^{\beta M_i I_i} \prod_{k \in \partial i \setminus j} \pi_{k \rightarrow i}(\sigma_k; \sigma_i), \\
\pi_i(\sigma_i) &\propto [\pi_i(\sigma_i)]^r \sum_{\{\sigma_k | k \in \partial i\}} e^{\beta M_i I_i} \prod_{k \in \partial i} \pi_{k \rightarrow i}(\sigma_k; \sigma_i),
\end{aligned} \tag{8}$$

where $r \geq 0$ is the reinforcement parameter. One can start with random initial values for the messages and update them according to the rBP equations in a random sequential way. At the beginning we set the reinforcement parameter to zero and increase its value slowly while the system converges to a solution. Notice that for $r \rightarrow \infty$ any solution of the problem is a fixed point of the above equations. The rBP equations thus suggest a local and distributive message passing algorithm to approach a Nash equilibrium.

B. 1RSB equations

For simplicity let us take the limit $\beta \rightarrow 0$ where $-\beta F$ is equivalent to entropy S . In the one-step replica symmetry breaking framework we assume there exist an exponentially large number of clusters of solutions [19]. This number is given by the so called complexity or configurational entropy by $e^{N\Sigma}$. Clusters can have different internal entropies (or sizes) $s \equiv S/N$ and $\Sigma(s)$ is used to indicate the complexity of a cluster of size s . The so called dominant clusters are those that maximize $\Sigma(s) + s$; with high probability a randomly selected solution belongs to these kind of clusters (though any specific algorithmic strategy would end up in clusters which are not necessarily the dominant ones). It is useful to

introduce Lagrange multiplier m and work with generalized free energy $m\Phi = \Sigma(s) + ms$. This allows to obtain the complexity by an inverse Legendre transform after computing the generalized free energy. As long as we are in the RS phase (even if the solution space is clustered) the relevant clusters are those corresponding to $m = 1$. This complexity goes continuously to zero at the thermodynamic RSB phase transition, after that the physical m would be less than 1, its value determined by the point of zero complexity.

We assume that each BP fixed point corresponds to a cluster or state of the system. The 1RSB equations give the statistics of BP messages among different clusters [19]

$$P_{i \rightarrow j}(\pi_{i \rightarrow j}) \propto \int \prod_{k \in \partial i \setminus j} dP_{k \rightarrow i}(\pi_{k \rightarrow i}) e^{m\Delta S_{i \rightarrow j}} \delta(\pi_{i \rightarrow j} - BP_{i \rightarrow j}), \quad (9)$$

where

$$e^{\Delta S_{i \rightarrow j}} = \sum_{\sigma_i, \sigma_{\partial i}} I_i \prod_{k \in \partial i \setminus j} \pi_{k \rightarrow i}(\sigma_k; \sigma_i). \quad (10)$$

Here m is called Parisi parameter and is to control the entropy density. Given the cavity distributions $P_{i \rightarrow j}(\pi_{i \rightarrow j})$ we obtain the 1RSB free energy in the Bethe approximation

$$\Phi = \sum_i \Delta\Phi_i - \sum_{(ij) \in \mathcal{G}} \Delta\Phi_{ij}, \quad (11)$$

where

$$e^{m\Delta\Phi_i} = \int \prod_{j \in \partial i} dP_{j \rightarrow i}(\pi_{j \rightarrow i}) e^{m\Delta S_i}, \quad (12)$$

$$e^{m\Delta\Phi_{ij}} = \int dP_{i \rightarrow j}(\pi_{i \rightarrow j}) dP_{j \rightarrow i}(\pi_{j \rightarrow i}) e^{m\Delta S_{ij}}.$$

The 1RSB free energy is related to the entropy and complexity by $m\Phi(m) = \Sigma(s) + ms$ where

$$m = -\frac{\partial \Sigma(s)}{\partial s}, \quad \Sigma(s) = -m^2 \frac{\partial \Phi(m)}{\partial m}. \quad (13)$$

The total 1RSB entropy is given by $\Sigma(s) + s$ computed at the physical m . As long as we are in the RS phase the 1RSB entropy computed at $m = 1$ is equal to the BP entropy.

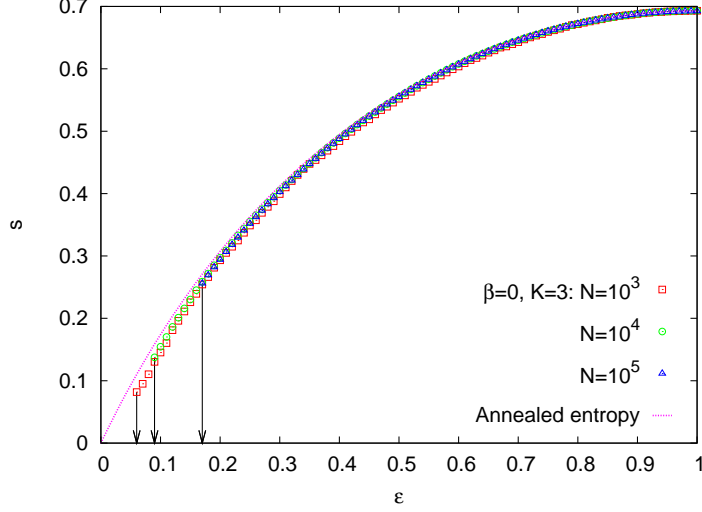


FIG. 1. BP entropy in the local graphical game with random payoffs.

For a given graphical game we can solve the 1RSB equations with the population dynamics technique [20, 22]. We represent the probabilities $P_{i \rightarrow j}(\pi_{i \rightarrow j})$ on each directed edge of the graph with a population $Pop_{i \rightarrow j}$ of BP messages $\{\pi_a | a = 1, \dots, N_p\}$. To update $Pop_{i \rightarrow j}$, we first select randomly messages π_{a_k} from $Pop_{k \rightarrow i}$ for $k \in \partial i \setminus j$. Then the new message $\pi_{i \rightarrow j}$ and the entropy shift $\Delta S_{i \rightarrow j}$ are computed according to the BP equations, and with probability $\propto e^{m \Delta S_{i \rightarrow j}}$ a randomly selected message from $Pop_{i \rightarrow j}$ is replaced with the new one. After converging to a stationary state we compute the average quantities by taking samples from the populations.

With the same scheme, one can write the 1RSB equations for other values of β , replacing the entropy with free energy.

C. Random payoffs

Let us start with fully random payoffs where each element $M_i(\sigma_i | \sigma_{\partial i}) \in [0, 1]$ is a uniform random number independent of the other payoffs. As figure 1 shows, here the BP entropy is positive only for approximate solutions with $\epsilon > \epsilon_c(N)$. For smaller ϵ the BP algorithm finds contradictory messages, something that usually happens when the solution set is empty. Moreover, the critical ϵ increases with N and finally for $N \rightarrow \infty$ we would have $\epsilon_c = 1$. Notice that for $\epsilon \geq 1$ any configuration is a solution to the problem.

It is easy to compute the average number of solutions or annealed entropy by averaging

over the randomness in the payoffs

$$s^{annealed} = \frac{1}{N} \ln \langle Z \rangle = \ln(2) + \ln \left(\frac{1 + 2\epsilon - \epsilon^2}{2} \right). \quad (14)$$

The above annealed entropy, displayed in figure 1, provides an upper bound for the correct number of solutions. However, this entropy is always greater than zero, except for exact solutions at $\epsilon = 0$. To understand why the approximate solutions do not survive in the thermodynamic limit we need to resort to another argument [41]: Consider an arbitrary region Ω of the graph \mathcal{G} . Suppose that we fix the strategies of boundary players $\partial\Omega$ to $\sigma_{\partial\Omega}$. Depending on the boundary state and payoffs, we may have no best response solution for the players in Ω . And this could happen for all boundary configurations $\sigma_{\partial\Omega}$. We denote the probability of this event by $P_{nosolution}(\Omega)$. Consider a collection \mathcal{C} of N_Ω disjoint regions in \mathcal{G} . Then, the probability of having a solution is overestimated by

$$P_{solution} \leq \prod_l (1 - P_{nosolution}(\Omega_l)). \quad (15)$$

The simplest choice is when each region consists of a single player. Obviously in this case $P_{nosolution} = 0$, which leads to a trivial inequality for $P_{solution}$. We can choose a larger region consisting of two neighboring players in the graph. For uniform and independent random payoffs one finds $P_{nosolution}(\Omega) \leq (1/8)^{2^{|\partial\Omega|}}$. As long as $|\partial\Omega|$ is finite and $N_\Omega = O(N)$ this results to an exponentially small probability of having a solution.

Any way, given a finite game, we find that BP always converges as long as $\epsilon > \epsilon_c(N)$; the system is in the RS phase with zero complexity and we can easily find an ϵ -Nash equilibrium using our reinforced BP equations. In figure 2 we compare the success probability of this algorithm with another message passing algorithm, similar to warning propagation, introduced in Ref. [34]. Here, a simple heuristic algorithm like Best Response (BR) does not converge to a solution. In the BR algorithm, we start from an initial strategy profile and as long as some players are not satisfied we randomly select one and update its strategy to the best response.

As mentioned above the random payoff ensemble has a trivial thermodynamic limit. To get around this problem we shall consider random games where nontrivial solutions exist.

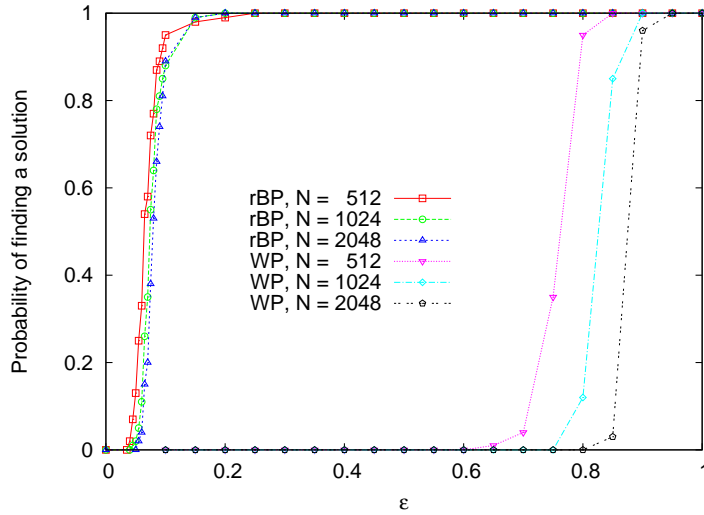


FIG. 2. Success probability of the rBP algorithm compared with the WP algorithm of Ref.[34] in the local graphical game with random payoffs. The data have been obtained by running the algorithms on 100 problem instances defined on random regular graphs of degree $K = 3$.

D. Random payoffs with hidden solutions

We can always modify random payoffs to ensure that our game has at least some pure Nash equilibria. Suppose that we want configuration $\underline{\sigma}^*$ to be a solution to the problem. Then we modify the payoffs in the following way: for each player i , if necessary, we swap the two values $M_i(\sigma_i^*|\sigma_{\partial i}^*)$ and $M_i(-\sigma_i^*|\sigma_{\partial i}^*)$ to satisfy the Nash condition for the player.

Let us take the uniform and independent random payoffs to see how the picture changes when we plant random configuration $\underline{\sigma}^*$ into the solution space. For each player we choose σ_i^* with equal probability from $\{-1, +1\}$. Figure 3 displays the BP entropy computed on some large instances of the problem. Interestingly, planting only one solution is enough to have an exponential number of pure Nash equilibria at $\epsilon = 0$. Moreover, these solutions do not disappear in the large N limit, as it happens for games with random payoff. To understand this consider the planted solution and a pair of neighboring players. There is a finite probability that after flipping the two strategies we get another Nash equilibrium. Moreover, there is an extensive number of independent such players that could result to an exponential number of solutions close to the planted one.

Introducing temperature into the problem makes the phase diagram more interesting in that we observe a critical line $\beta_c(\epsilon)$ separating the RS and RSB phases in the $\beta - \epsilon$ plane,

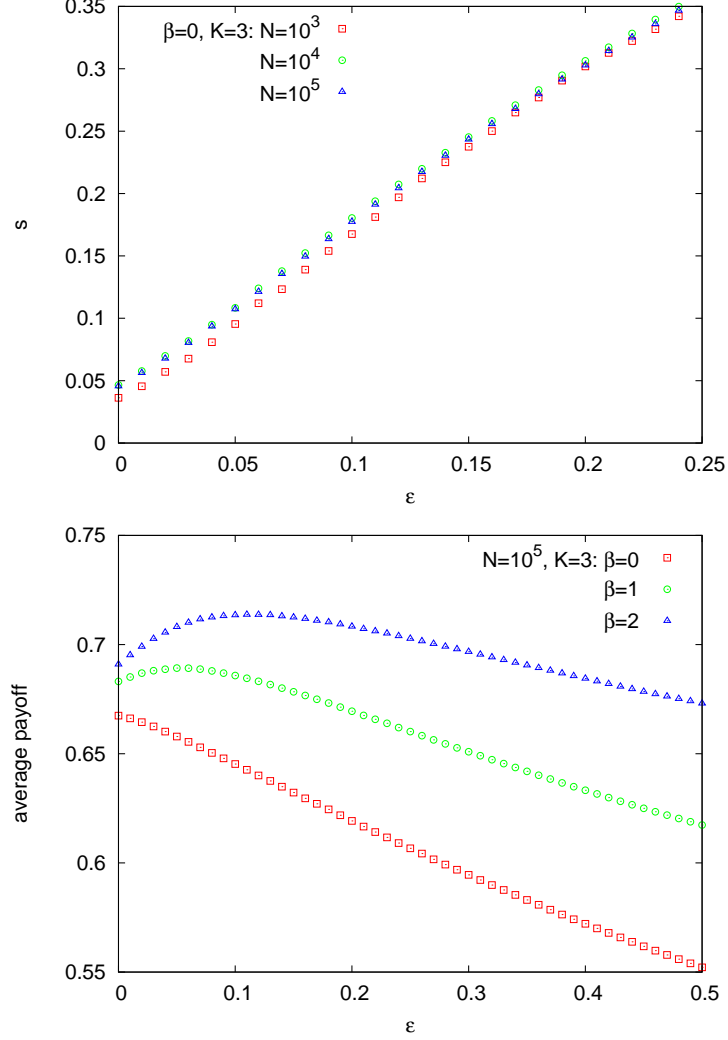


FIG. 3. BP entropy and average payoff in the local graphical game with a hidden solution.

see figure 4; BP does not converge for $\beta > \beta_c$, signaling an RSB phase transition. Moreover, there is a finite temperature-gap for any ϵ , meaning that finding a typical ϵ -Nash equilibrium is an easy task. Figure 5 shows how the entropy and average payoff change with β for exact Nash equilibria. For comparison we have also given the average payoff of solutions obtained with the rBP algorithm. In figure 3 we also display the average total payoff as a function of ϵ for different values of β . We observe a maximum appearing in the average payoff as we increase β ; obviously for small β , more accurate solutions have larger total payoff than solutions satisfying Nash condition at a larger ϵ , since Nash condition is already maximizing the local payoffs. However, when β is large enough we select those strategies maximizing the total payoff and with a larger ϵ we have more space to find a better global maximum.

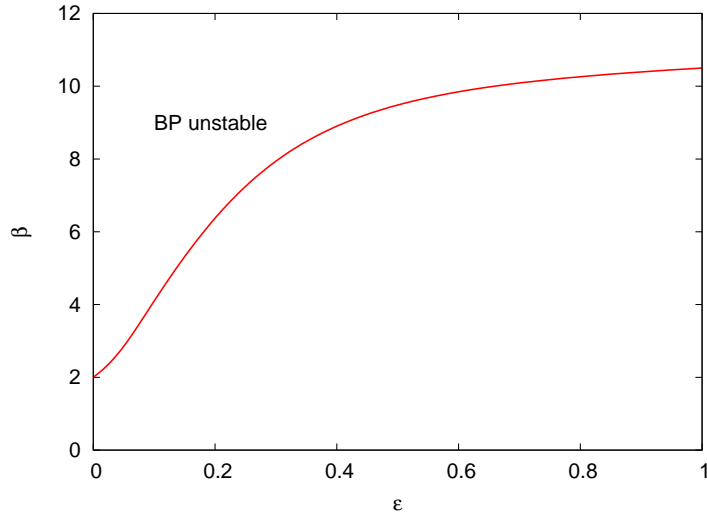


FIG. 4. The phase diagram in the local graphical game with a hidden solution. The line has been obtained by checking the convergence of BP equations on a single instance of size $N = 10^5$ and degree $K = 3$. For $\beta > \beta_c$ the equations do not converge in $T_{max} = 1000$ iterations.

A more accurate estimate of the entropy and of the total payoff is obtained by considering replica symmetry breaking. Figure 6 shows the $m = 1$ complexity computed in the 1RSB approximation at different temperatures. The system is in the RS phase for small β ; BP converges, complexity is zero and the total 1RSB entropy is equal to the BP entropy. For larger β we have replica symmetry breaking, more precisely we enter into a condensed phase where only a finite number of solution clusters are relevant. In this case, as we see in the figure, the $m = 1$ complexity is negative but the relevant m will be less than 1 where complexity is zero.

Another way of planting solutions is to modify the random payoffs to mimic the constraints in an already known problem. In this case, we go through all the neighboring configurations of a player and if necessary swap the two values $M_i(\sigma_i|\sigma_{\partial i})$ and $M_i(-\sigma_i|\sigma_{\partial i})$ according to the constraints. An example that we will later study in this paper is the maximal independent set (mIS) problem where a player plays +1 only if all its neighbors play -1. Regarding the payoffs, it means that for each player we need to have $M_i(+1|\forall j \in \partial i, \sigma_j = -1) \geq M_i(-1|\forall j \in \partial i, \sigma_j = -1)$ and $M_i(-1|\exists j \in \partial i, \sigma_j = +1) \geq M_i(+1|\exists j \in \partial i, \sigma_j = +1)$. It is easy to obtain a typical solution for this problem by running the Best Response algorithm. What is difficult is to find an optimal solution, for example with a large number of active

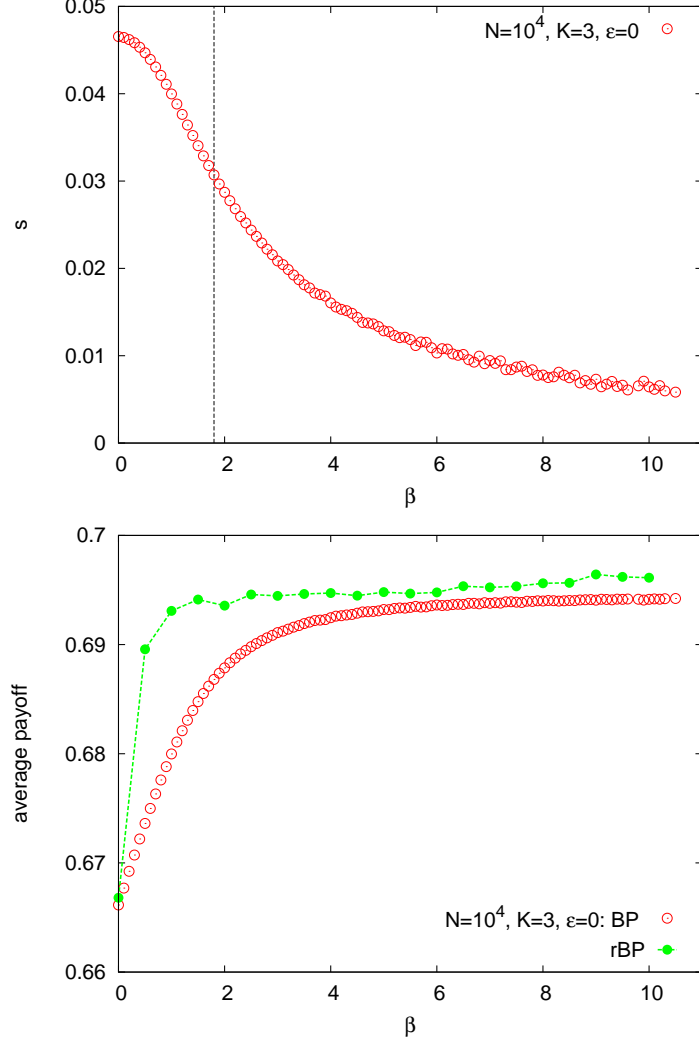


FIG. 5. BP entropy and average payoff in the local graphical game with a hidden solution. The vertical line separates the RS (small β) and RSB (large β) regions. The rBP average payoffs are results of averaging over at least 50 solutions obtained with the rBP algorithm.

players. Algorithms based on the BP equations help us to find good optimal solutions in large problem instances [42]. In figure 7 we show the entropy of solutions with magnetization $m = \frac{1}{N}(\sum_i \sigma_i)$ in the region that BP equations converge.

IV. GLOBAL GRAPHICAL GAMES

In a global graphical game the payoffs depend, besides the local neighborhood, on the global state of the system, for instance as

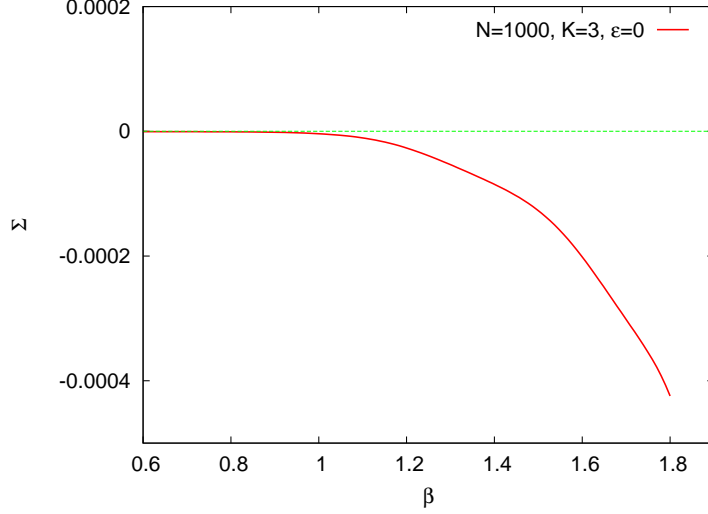


FIG. 6. $m = 1$ complexity in the local graphical game with a hidden solution. The data are obtained by solving 1RSB equations with population dynamics on a single instance of the problem. On each directed edge of the graph we have a population of size 1000.

$$M_i = M_i^{local}(\sigma_i|\sigma_{\partial i}) + M_i^{global}(\sigma_i|g), \quad (16)$$

where $g(\underline{\sigma})$ is an aggregate quantity depending on the strategy profile. As for local games the total number of solutions can be written as

$$Z = \sum_{\underline{\sigma}} \prod_i I_i(\sigma_i|\sigma_{\partial i}, g), \quad (17)$$

where I_i is an indicator function to check Nash condition for player i . An interesting example is

$$M_i^{global} = h\sigma_i m, \quad m = \frac{1}{N} \sum_i \sigma_i, \quad (18)$$

where players will receive more than their local payoffs if they are in majority or minority, depending on the sign of global field h . These kind of interactions are similar to mean field models in statistical physics but in a different setting. It also makes sense from a social point of view to have such global incentives [43].

The sole global problem $M_i = h\sigma_i m$ has only two solutions with magnetizations $m = \pm 1$ for $h > 0$ and $C(N, N/2)$ solutions with magnetization $m = 0$ when $h < 0$ and N is even.

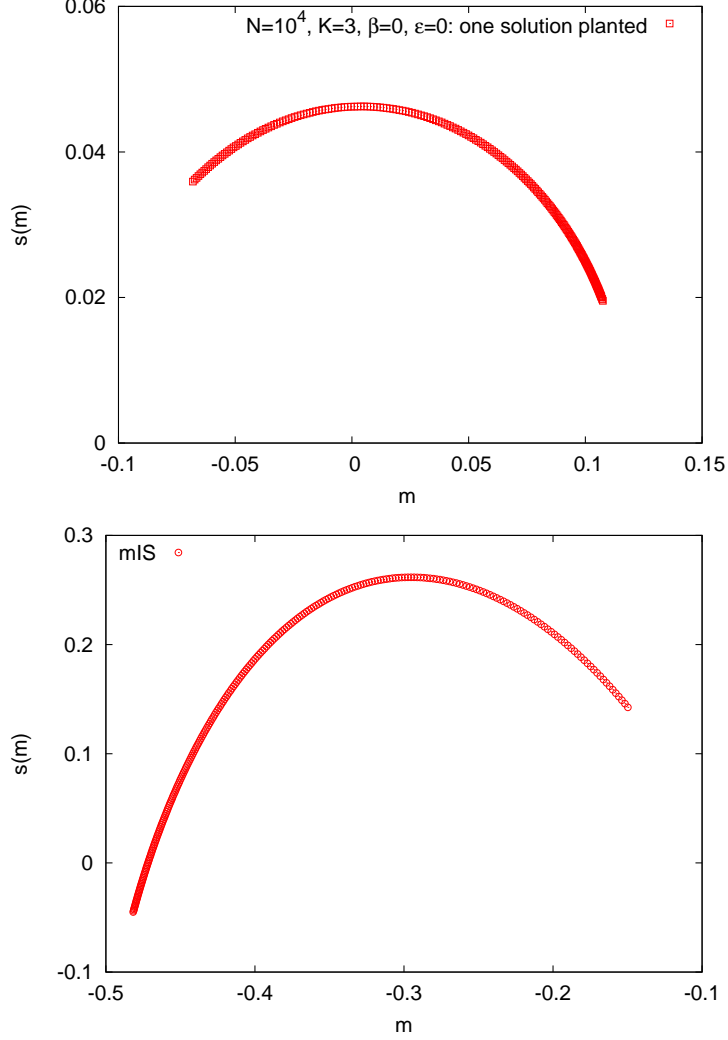


FIG. 7. Entropy as a function of magnetization in the local graphical game with hidden solutions. We plot the entropy only if the BP fixed point is stable.

Here $C(N, l) = \binom{N}{l}$ is the Binomial coefficient.

One could consider other global terms more suited to the local problem. For instance, we might be interested on the total parity of solutions, where payoffs are given by

$$M_i = M_i^{local}(\sigma_i | \sigma_{\partial i}) + h \prod_i \sigma_i. \quad (19)$$

In statistical physics these type of interactions are studied to model structural glasses. In this case the global problem $M_i = h \prod_i \sigma_i$ partitions the configuration space to even and odd parity solutions for $h > 0$ and $h < 0$, respectively.

In this paper we shall study the former problem where total activity or magnetization m determines the global payoff. For the sake of simplicity, in the rest of the paper we shall work in the replica symmetric approximation and set $\beta = 0$ and $\epsilon = 0$.

A. Rigorous results

Suppose that local payoffs take integer values in $\{0, 1\}$ and there is no degeneracy, that is $M_i^{local}(\sigma_i|\sigma_{\partial i}) \neq M_i^{local}(-\sigma_i|\sigma_{\partial i})$. This special case helps us to understand the problem better when studying real and random payoffs. Consider a strategy profile $\underline{\sigma}$ with magnetization m . Flipping the strategy σ_i of player i results to the following change in her payoff:

$$\Delta M_i = \Delta M_i^{local} - 2h(\sigma_i m - \frac{1}{N}), \quad (20)$$

where the $1/N$ term comes from the change in total magnetization. In order to simplify the arguments, we will consider only strict solutions, i.e. those with $\Delta M_i < 0$. We say a player is locally happy if $\Delta M_i^{local} < 0$ and group the players in distinct sets H_+, H_-, U_+, U_- for locally happy (H) and unhappy (U) players with plus and minus strategies. It is easy to see that players in H_+ satisfy the Nash condition only if:

$$\begin{cases} m - \frac{1}{N} < \frac{1}{2|h|}, & h < 0; \\ m - \frac{1}{N} > -\frac{1}{2|h|}, & h > 0; \end{cases} \quad (21)$$

The above conditions define two lines in the $m - h$ plane. Similarly we can write the conditions for the other sets. From these we conclude that:

Remark (1): For $-\frac{1}{2(1-1/N)} < h < \frac{1}{2(1+1/N)}$ any local solution is also a solution of the problem. A local solution is a solution of type (H_+, H_-) . The entropy of local solutions survived at global field h is:

$$s^{local}(h) = \max_{-m^*(h) < m < +m^*(h)} s^{local}(m), \quad (22)$$

where

$$m^*(h) = \min \left(1, \frac{1}{2|h|} - \frac{\text{sgn}(h)}{N} \right). \quad (23)$$

We defined $s^{local}(m)$ as the entropy of local solutions (i.e. at $h = 0$) having magnetization m .

Remark (2): For $h < -N/2$ and $h > N/2$ only global solutions remain. For positive h the two global solutions with $m = \pm 1$ appear at $h = \frac{1}{2(1-1/N)}$, whereas for negative h the global solutions with $m = 0$ appear at $h = -N/2$.

Remark (3): For $-N/2 < h < N/2$ one may have mixed solutions. These are solutions of type (H_+, H_-, U_+) (for $hm > 0$) or (H_+, H_-, U_-) (for $hm < 0$). That is we can not have solutions with both sets U_+ and U_- non-empty. There is no mixed solution for $h > 0$, and if there is a mixed solution for $h < -1/2$ it must have magnetization $m = \pm \frac{1}{2|h|}$.

When the local payoffs are uniform random numbers in $[0, 1]$ satisfying the maximal independent set constraints, we have:

Remark (4): There is no solution with magnetization $m \in]2\rho_{max} - 1, 0[$ for $h > 0$, where ρ_{max} denotes the size of the maximum independent set in graph \mathcal{G} . Moreover, the probability of having a solution is zero in the thermodynamic limit when $0 < 2h(m \pm 1/N) < 1$.

The reader can find more about these remarks in appendix A.

B. Annealed approximation

Given the ensemble of local payoffs we can compute the average number of solutions in a graphical game as

$$\langle Z \rangle = \langle e^{Ns} \rangle = \sum_{\underline{\sigma}} \prod_i \langle I_i \rangle. \quad (24)$$

The average is taken over the randomness in the independent local payoffs. The convexity of exponential function ensures that $s^{annealed} \equiv \ln \langle Z \rangle$ is an overestimate of the average entropy $\langle \ln Z \rangle$. In a global graphical game the constraint I_i depends also on total magnetization density m , therefore, we will restrict the above equations to the subspace of fixed $M = \sum_i \sigma_i$. Then the average number of solutions reads

$$\langle Z \rangle = \sum_M \sum_{N(H_+), N(H_-)} e^{Ns[m, n(H_+), n(H_-)]} \prod_{\sigma=+1, -1} p(H_\sigma)^{N(H_\sigma)} p(U_\sigma)^{N(U_\sigma)}, \quad (25)$$

where $e^{Ns[m, n(H_+), n(H_-)]}$ is the number of configurations with specified densities, e.g. $n(H_+) = N(H_+)/N$. Notice that the other two densities are not independent but given by

$$\begin{aligned} N(U_+) &= \frac{N+M}{2} - N(H_+), \\ N(U_-) &= \frac{N-M}{2} - N(H_-). \end{aligned} \quad (26)$$

Moreover, as described in the previous section we can not have all densities nonzero; depending on the sign of hm one of the two quantities $N(U_+), N(U_-)$ should be zero, such that at the end there remains only one independent parameter $N(H_+)$ or $N(H_-)$. For given h and m , the probability that a locally happy player is satisfied by the total payoff is

$$\begin{aligned} p(H_-) &= \Pr(\Delta M^{local} > 2h(m + 1/N)), \\ p(H_+) &= \Pr(\Delta M^{local} > -2h(m - 1/N)). \end{aligned} \quad (27)$$

And if the player is locally unhappy

$$\begin{aligned} p(U_-) &= \Pr(\Delta M^{local} < -2h(m + 1/N)), \\ p(U_+) &= \Pr(\Delta M^{local} < 2h(m - 1/N)). \end{aligned} \quad (28)$$

For uniform random payoffs in $[0, 1]$ we have

$$\Pr(|\Delta M^{local}| > \epsilon) = \begin{cases} 1 & \epsilon < 0; \\ (1 - \epsilon)^2 & 0 \leq \epsilon \leq 1; \\ 0 & \epsilon > 1. \end{cases} \quad (29)$$

The above quantities are enough to compute the average number of solutions for random payoffs. Let us separate the two cases of positive and negative fields. For $h > 0$:

$$\langle Z \rangle = \frac{1}{2^N} \begin{cases} C(N, (N-M)/2) p(H_+)^{(N+M)/2} [1 + p(U_-)]^{(N-M)/2} & M < -1; \\ C(N, (N-M)/2) p(H_+)^{(N+M)/2} & M = -1; \\ C(N, N/2) p(H_-)^{N/2} p(H_+)^{N/2} & M = 0; \\ C(N, (N+M)/2) p(H_-)^{(N-M)/2} & M = +1; \\ C(N, (N+M)/2) p(H_-)^{(N-M)/2} [1 + p(U_+)]^{(N+M)/2} & M > +1. \end{cases} \quad (30)$$

And for $h < 0$:

$$\langle Z \rangle = \frac{1}{2^N} \begin{cases} C(N, (N+M)/2) p(H_-)^{(N-M)/2} [1 + p(U_+)]^{(N+M)/2} & M < 0; \\ C(N, N/2) [1 + p(U_-)]^{N/2} [1 + p(U_+)]^{N/2} & M = 0; \\ C(N, (N-M)/2) p(H_+)^{(N+M)/2} [1 + p(U_-)]^{(N-M)/2} & M > 0; \end{cases} \quad (31)$$

Notice that given h the entropy is symmetric with respect to m . It is easy to see that the only positive contribution to the entropy comes from $m = 0$ configurations when $h < 0$ and scales with N :

$$s^{annealed} = \ln \left[2 - \left(1 - \frac{2|h|}{N} \right)^2 \right]. \quad (32)$$

In the other cases, the annealed entropy as a function of magnetization is always less than or equal to zero. Consider for example the case $h > 0$ such that $0 < 2hm < 1$, we have

$$\begin{aligned} s^{annealed} = & -\ln(2) - \left(\frac{1+m}{2} \right) \ln \left(\frac{1+m}{2} \right) - \left(\frac{1-m}{2} \right) \ln \left(\frac{1-m}{2} \right) \\ & + \left(\frac{1+m}{2} \right) \ln[1 + 4hm - 4h^2m^2] + \left(\frac{1-m}{2} \right) \ln[1 - 4hm + 4h^2m^2]. \end{aligned} \quad (33)$$

The above entropy is zero only at $m = 0$, or at $m = m_0(h)$ if h is greater than critical value $h_c \approx 0.25$ determined by the following equation

$$\begin{aligned} m_0 &= \tanh(h_g - h_l), \\ h_l &= 2h(1 - m_0) \frac{1 - 2hm_0}{1 - 4hm_0 + 4h^2m_0^2} + \frac{1}{2} \ln[1 - 4hm_0 + 4h^2m_0^2], \\ h_g &= 2h(1 + m_0) \frac{1 - 2hm_0}{1 + 4hm_0 - 4h^2m_0^2} + \frac{1}{2} \ln[1 + 4hm_0 - 4h^2m_0^2]. \end{aligned} \quad (34)$$

The nontrivial magnetization m_0 approaches to 1 as global field h reaches the value $1/2$.

The situation is a bit more complex when the local payoffs have some structure. The difficulty comes from computing $s[m, n(H_+), n(H_-)]$ which was trivial for random payoffs. When dependency graph \mathcal{G} is a chain we can compute this entropy exactly. However, for arbitrary graphs we will estimate it in the Bethe approximation by solving the following problem:

$$Z(x, y) = \sum_{\underline{\sigma}} e^{xN(H_+) + yN(H_-)} I(U_\tau = \emptyset). \quad (35)$$

Here $\tau = -\text{sgn}(hm)$ and $I(U_\tau = \emptyset)$ is an indicator function to have U_+ or U_- empty, depending on the sign of hm . The entropy is obtained by a Legendre transformation after computing the free energy of this problem for appropriate values of fields x and y . Figure 8 shows the annealed entropy obtained in this way for random regular graphs of degree $K = 3$ with local mIS constraints. Notice the small entropy maximums appearing close to the global polarized solutions when h is approaching $h_c \approx 0.43$. We remind that according to remark (4) the entropy is zero in thermodynamic limit for $0 < 2h(m \pm 1/N) < 1$. As for random local games in section III C, here the annealed entropy does not give the correct behavior. However, it is still useful in that we obtain a qualitative picture of the entropy in finite size systems.

C. Global algorithms

The global graphical game can be treated like the local one by introducing a global constraint fixing the global quantity. Suppose that the global payoff depends on a global variable $g \in \Lambda_g$. We write the following partition function to count the Nash equilibria

$$Z = \sum_g \sum_{\underline{\sigma}} I_g(\underline{\sigma}) \prod_i I_i(\sigma_i | \sigma_{\partial i}, g), \quad (36)$$

where as before I_i checks for Nash condition and I_g is an indicator function to fix quantity $g = g(\underline{\sigma})$. We can then write the standard BP equations regarding I_g as another constraint in the problem. However, in this way we introduce a large number of loops into the problem, which destroys the BP exactness even when the original graph \mathcal{G} is a tree. To preserve this property here we follow another strategy in which the global constraint is broken to local ones in the expense of introducing new variables. Consider global quantities that their computation can be partitioned into smaller local computations. This is the case for example when g is total magnetization or parity. Then we introduce cavity variables $g_{i \rightarrow j}$ which are passed along a spanning tree \mathbb{T} of graph \mathcal{G} . These variables are determined by other cavity

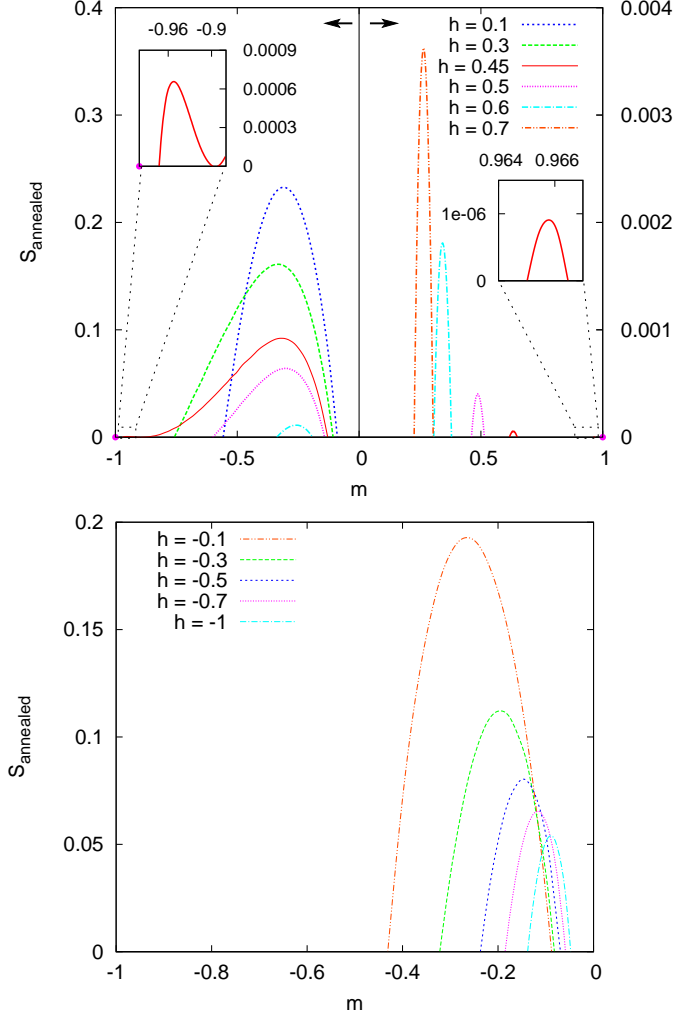


FIG. 8. Annealed entropy as a function of magnetization in the global problem with local mIS constraints. The entropy has been computed for random regular graphs of degree $K = 3$.

variables as $g_{i \rightarrow j} = g(\sigma_i, \{g_{k \rightarrow i} | k \in \partial i \setminus j, \mathbb{T}\})$. For instance, when g is total magnetization, the cavity magnetizations are given by $g_{i \rightarrow j} = \sigma_i + \sum_{k \in \partial i \setminus j, \mathbb{T}} g_{k \rightarrow i}$. Now the global partition function can be rewritten as

$$Z = \sum_{\{g_{i \rightarrow j} | (ij) \in \mathbb{T}\}} \sum_{\underline{\sigma}} \prod_i I_{ig}(\sigma_i | \sigma_{\partial i}, g_i), \quad (37)$$

where I_{ig} is to check the Nash condition and outgoing cavity variables $\{g_{i \rightarrow j} | j \in \partial i, \mathbb{T}\}$. Notice that each player computes its estimate of the global quantity $g_i = g(\sigma_i, \{g_{k \rightarrow i} | k \in \partial i, \mathbb{T}\})$ locally after receiving the incoming cavity variables. The BP equations with new variables read

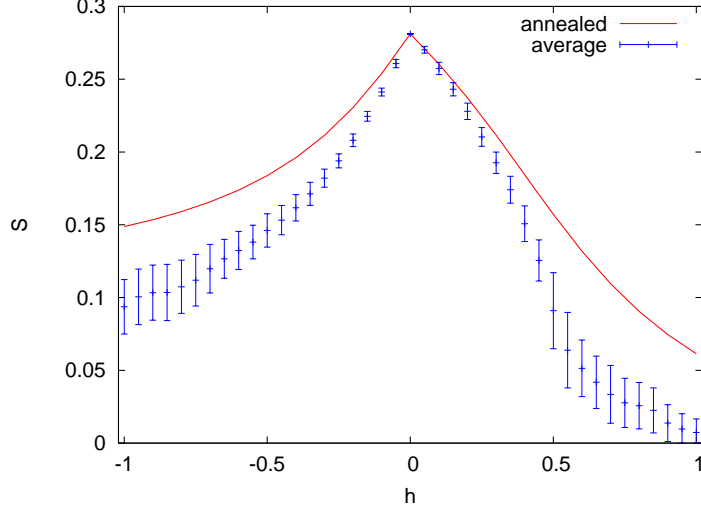


FIG. 9. Comparing the exact entropy with the annealed one for a chain of players in the global problem with local mIS constraints. The data are results of averaging over 10 instances of random payoffs.

$$\pi_{i \rightarrow j}(\sigma_i, g_{i \rightarrow j}; \sigma_j, g_{j \rightarrow i}) \propto \sum_{\{\sigma_k, g_{k \rightarrow i} | k \in \partial i \setminus j\}} I_{ig}(\sigma_i | \sigma_{\partial i}, g_i) \prod_{k \in \partial i \setminus j} \pi_{k \rightarrow i}(\sigma_k, g_{k \rightarrow i}; \sigma_i, g_{i \rightarrow k}). \quad (38)$$

The time complexity of this algorithm is $NK|\Lambda_g|^{k_{max}}$ where k_{max} is the maximum degree in spanning tree T . Actually, the complexity can be reduced to $NK|\Lambda_g|^2$ if we pass the messages $g_{i \rightarrow j}$ along a spanning chain which is used just to compute the global quantity. Nevertheless, the algorithm is still computationally expensive specially when $|\Lambda_g|$ is large.

In figure 9 we compare the exact BP entropy computed in this way with the annealed entropy in a small chain of players, when $g = m$ and $|\Lambda_g| = N$.

D. Local algorithms for global games

In this section we present another way of dealing with the global payoff when it depends on the total magnetization. Considering h and m as fixed parameters in the payoffs, we reduce the problem to a local one but with modified conditions for solutions. The partition function for this local problem reads

$$Z(m) = \sum_{\underline{\sigma}} \prod_i I_i(\sigma_i | \sigma_{\partial i}, m). \quad (39)$$

Clearly the entropy computed in this way is an upper bound for $s(m)$, the entropy of solutions with magnetization m in the global problem. The reason is that here m is just a parameter which is not necessarily the total magnetization. Indeed, we can do better than this by introducing an external field to really fix the total magnetization to m

$$Z(m) = \sum_{\underline{\sigma}} e^{x \sum_i \sigma_i} \prod_i I_i(\sigma_i | \sigma_{\partial i}, m), \quad (40)$$

where x is chosen such that

$$m = \frac{1}{N} \frac{\partial \ln Z(m)}{\partial x}. \quad (41)$$

In the Bethe approximation the average magnetization is computed from the cavity messages satisfying BP equations

$$\pi_{i \rightarrow j}(\sigma_i; \sigma_j) \propto \sum_{\{\sigma_k | k \in \partial i \setminus j\}} e^{x \sum_i \sigma_i} I_i(\sigma_i | \sigma_{\partial i}, m) \prod_{k \in \partial i \setminus j} \pi_{k \rightarrow i}(\sigma_k; \sigma_i). \quad (42)$$

In this way we obtain a better estimation of $s(m)$ as displayed in figure 10. This method of fixing magnetization has already been implemented in Refs.[44–46].

The above entropy can also be computed with the population dynamics technique. To do this we represent the set of BP messages in a graph by a large population of messages. The population is updated by selecting randomly $K - 1$ messages π_{a_i} and computing a new BP messages according to the BP equations. Then a randomly selected element of population is replaced with the new message. To fix the magnetization we also update field x such that the expected magnetization in population is equal to m .

There are some subtle points here to mention about the population dynamics. First note that in the population dynamics we do not have the small payoff change $2h/N$ in the Nash conditions; we are at the thermodynamic limit $N \rightarrow \infty$, and h is finite. In all the numerical simulations we will work with a global field of order 1, which is reasonable since the local payoffs are real numbers in $[0, 1]$. At the first sight the small term $2h/N$ seems irrelevant for large problem sizes but adding this correction could eliminate some solutions when h is

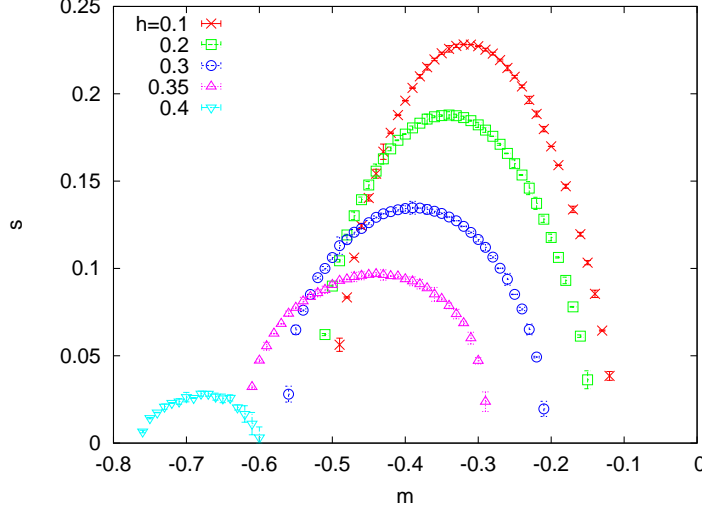


FIG. 10. BP entropy as a function of magnetization in the global problem with local mIS constraints. The data have been obtained with population dynamics with $N_p = 2 \times 10^5$. Here the errorbars show the difference between the fixed and converged magnetizations.

positive. We recall that the Nash condition for player i is $\Delta M_i^{local} - 2h\sigma_i m \leq -2h/N$; and this is a stronger condition for a solution than $\Delta M_i^{local} - 2h\sigma_i m \leq 0$ when $h > 0$. Indeed, there is a finite probability to miss a solution by adding the small term $2h/N$ even in the large N limit. Secondly, the size of population that we use to represent the statistics of BP messages is finite. It means that, even if in the thermodynamic limit we do not have any solution, we may still observe a positive entropy due to the finite size of the population. What we can do is just to work with the largest population allowed in our numerical simulations.

The local problem can also be used to converge the system to a problem solution by applying reinforcement. Here are the reinforced BP equations

$$\begin{aligned} \pi_{i \rightarrow j}(\sigma_i; \sigma_j) &\propto [\pi_i(\sigma_i)]^r \sum_{\{\sigma_k | k \in \partial i \setminus j\}} I_i(\sigma_i | \sigma_{\partial i}, m) \prod_{k \in \partial i \setminus j} \pi_{k \rightarrow i}(\sigma_k; \sigma_i), \\ \pi_i(\sigma_i) &\propto [\pi_i(\sigma_i)]^r \sum_{\{\sigma_k | k \in \partial i\}} I_i(\sigma_i | \sigma_{\partial i}, m) \prod_{k \in \partial i} \pi_{k \rightarrow i}(\sigma_k; \sigma_i). \end{aligned} \quad (43)$$

The algorithm works by computing $m = \frac{1}{N} \sum_i [\pi_i(+1) - \pi_i(-1)]$ from the BP marginals at each iteration and using that as an estimation of total magnetization in the Nash conditions. Remember that m can be computed locally by passing appropriate messages $g_{i \rightarrow j}$ along a spanning tree. Figure 11 compares the magnetization of solutions found in this way with the

Best Response algorithm in a problem with local mIS constraints. We see that even for a large graph of $N = 10^4$ players, there are still some Nash equilibria for $0 < 2hm < 1$, where the solution set is asymptotically empty. The Best Response algorithm finds easily a solution when the global field is positive but it does not converge for negative h . Using the above rBP algorithm we could obtain different kinds of solutions depending on the reinforcement parameter. The solutions lay, of course, in the region that BP equations converge. In figure 11 we also show the typical entropy for positive values of the global field. It seems that both the entropy and magnetization approach continuously to their limiting global values.

V. CONCLUSION

We used the cavity method in conjunction with rigorous bounds to study random graphical games with local and global interactions. We analyzed the phase diagram of the problem and presented some local message passing algorithms which allow to find efficiently a Nash equilibrium.

More specifically, we studied graphical games with local payoffs coming from the maximal independent set problem and global payoffs which depend on the average strategy over the whole graph (the so called total magnetization in the physics jargon). Introducing the global interaction resulted to a new set of equilibria which are a mixture of locally happy and unhappy players. In summary: (i) Using rigorous arguments and the annealed entropy, we showed that these equilibria cannot be present in some regions of the phase space in the thermodynamic limit. (ii) We observed an exponentially large number of these equilibria for positive global fields and negative magnetizations which we conjecture survive in the thermodynamic limit. Indeed, a simple heuristic algorithm like the Best Response is able to find such a typical solution in very large problem instances. The entropy and the total magnetization of the typical equilibria decrease continuously as we increase the global field starting from zero. (iii) Numerical simulations and annealed approximation results show the existence of a critical value of h above which a cluster of solutions dominated by the global interaction appears. Close to this value of the global field, the typical entropy and magnetization decrease very rapidly, separating the regions governed by the local and global interactions.

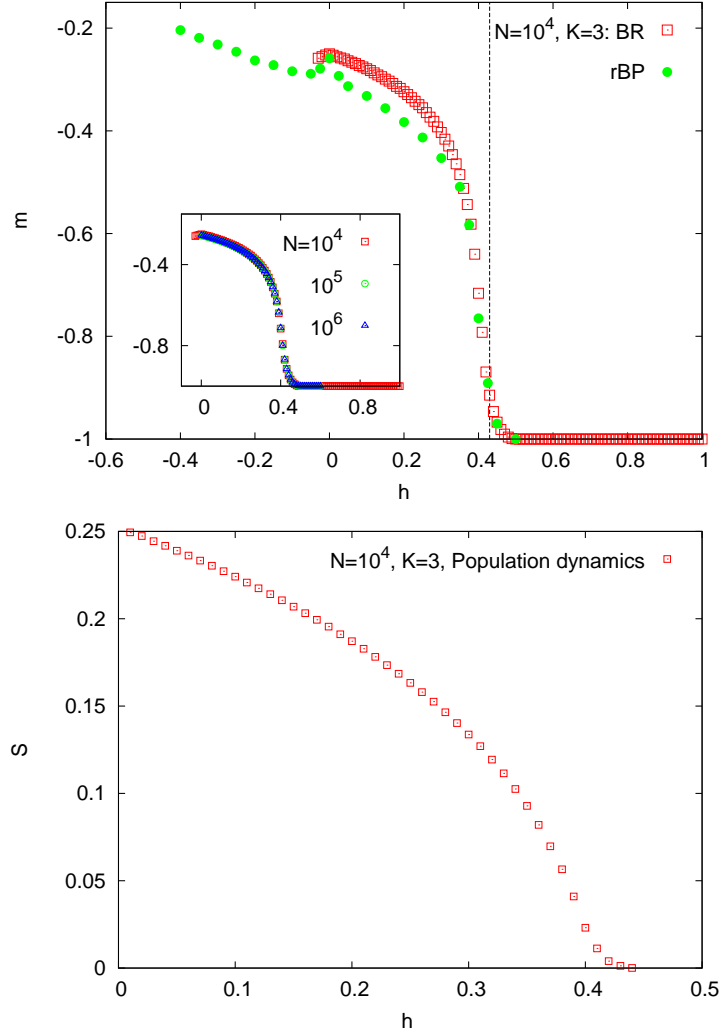


FIG. 11. Upper plot: average magnetization of solutions found by the rBP and BR algorithms in 100 instances of the global problem with local mIS constraints. The vertical line shows the point where globally dominated solutions appear. The inset displays the behavior for different sizes. Lower plot: typical entropy computed with population dynamics.

ACKNOWLEDGMENTS

We would like to thank A. Braunstein, L. Dall'Asta and M. Marsili for useful discussions.

Appendix A: More details on Remarks (1) to (4)

On Remark (1): Consider a solution of the problem with magnetization m and global field h . In a local solution players are either in set H_+ or H_- . For a plus player the Nash

condition is $\Delta M^{local} - 2h(m - 1/N) < 0$ which for happy players means $-1 - 2h(m - 1/N) < 0$. Notice that for integer payoffs in $\{0, 1\}$ a happy player has $\Delta M^{local} = -1$ whereas for an unhappy player $\Delta M^{local} = +1$. Therefore, a happy plus player is plying the best response as long as

$$\begin{cases} m - \frac{1}{N} < \frac{1}{2|h|}, & h < 0; \\ m - \frac{1}{N} > -\frac{1}{2|h|}, & h > 0; \end{cases} \quad (\text{A1})$$

For a happy minus player the Nash condition is $-1 + 2h(m + 1/N) < 0$ which is satisfied when

$$\begin{cases} m + \frac{1}{N} > -\frac{1}{2|h|}, & h < 0; \\ m + \frac{1}{N} < \frac{1}{2|h|}, & h > 0; \end{cases} \quad (\text{A2})$$

To have a local solution we just need $-m^*(h) < m < m^*(h)$ where $m^*(h) = \min(1, \frac{1}{2|h|} - \frac{\text{sgn}(h)}{N})$. If $|h| < \frac{1}{2(1+\text{sgn}(h)/N)}$ we get $m^* = 1$, that is all local solutions survive after adding the global incentives.

On Remark (2): According to the above arguments, to have both happy plus and minus players when $h > N/2$ we need at the same time $m > 0$ and $m < 0$, which is impossible. But we can have all plus or all minus solutions, that is the two global solutions for $h > 0$. Indeed these two solutions appear at $h = \frac{1}{2(1-1/N)}$ where unhappy players are allowed. To see this we note that a player in U_+ is satisfied if

$$\begin{cases} m - \frac{1}{N} < -\frac{1}{2|h|}, & h < 0; \\ m - \frac{1}{N} > \frac{1}{2|h|}, & h > 0; \end{cases} \quad (\text{A3})$$

The all plus solution is possible if $1 > \frac{1}{2h} + \frac{1}{N}$, that is $h > \frac{1}{2} + \frac{h}{N}$. For players in U_- we have

$$\begin{cases} m + \frac{1}{N} > \frac{1}{2|h|}, & h < 0; \\ m + \frac{1}{N} < -\frac{1}{2|h|}, & h > 0; \end{cases} \quad (\text{A4})$$

The all minus solution is possible if $-1 < -\frac{1}{2h} - \frac{1}{N}$, that is $h > \frac{1}{2(1-1/N)}$.

On the other hand, to have the $m = 0$ global solutions for $h < 0$, we need $0 > \frac{1}{2|h|} - \frac{1}{N}$ which implies $|h| > N/2$.

On Remark (3): From the above equations we see that to have players in both U_+ and U_- when $h > 0$, we need $m > \frac{1}{2|h|} + \frac{1}{N}$ and $m < -\frac{1}{2|h|} - \frac{1}{N}$ which is not possible. Indeed one could have only unhappy plus players if $m > 0$ or unhappy minus players if $m < 0$. But a mixed solution should contain both plus and minus players, otherwise it would be completely polarized global solution. Considering the case $m > 0$, the conditions for unhappy plus players and happy minus players give contradictory inequalities $m > \frac{1}{2|h|} + \frac{1}{N}$ and $m < \frac{1}{2|h|} - \frac{1}{N}$. In the case $m < 0$, the conditions for unhappy minus players and happy plus players give contradictory inequalities $m > -\frac{1}{2|h|} + \frac{1}{N}$ and $m < -\frac{1}{2|h|} - \frac{1}{N}$. Therefore, we can not have a mixed solution for positive global fields.

When $h < 0$, to have non-empty sets U_+ and U_- we need $m < -\frac{1}{2|h|} + \frac{1}{N}$ and $m > \frac{1}{2|h|} - \frac{1}{N}$. Again we can only unhappy plus players if $m < 0$ or unhappy minus players if $m > 0$. In other words, we can only have solutions of type (H_+, H_-, U_+) if $hm > 0$ and (H_+, H_-, U_-) if $hm < 0$. Considering the two cases $m > 0$ and $m < 0$ we find $\frac{1}{2|h|} - \frac{1}{N} < m < \frac{1}{2|h|} + \frac{1}{N}$ and $-\frac{1}{2|h|} - \frac{1}{N} < m < -\frac{1}{2|h|} + \frac{1}{N}$, respectively. That is, if there is a mixed solution it should have magnetization $\pm \frac{1}{2|h|}$ and so $2|h| > 1$.

Note also that for both $h > 0$ and $h < 0$ we can not have an $m = 0$ solution of type (H_+, H_-, U_+, U_-) as long as $|h| < N/2$.

On Remark (4): Consider the maximal independent set problem on graph \mathcal{G} and let ρ_{max} be the size of a maximum independent set. Thus, all local solutions have a magnetization less than $2\rho_{max} - 1$. Consider the global problem when $h > 0$. According to the previous remarks, if there is a solution of magnetization $2\rho_{max} - 1 < m < 0$ to the problem, it should be of type (H_+, H_-, U_-) . But this is not possible because we can always change the state of a player in U_- to $+1$, increasing the magnetization and still respecting the mIS constraints. In other words, such a configuration would lead to local solution of magnetization $m > 2\rho_{max} - 1$, which is of course a contradiction.

The other part of remark (4) can be proved by bounding the probability of having a solution, as we did in section IIIC for random payoffs. Consider a region Ω which is a chain of three neighboring players (j, i, k) . Here we show that if $0 < 2h(m \pm 1/N) < 1$, then for any boundary configuration $\sigma_{\partial\Omega}$ there is a nonzero probability of having no solution in the region.

Let us start by eliminating the two configurations $(+1, +1, -1), (-1, +1, +1)$ in which player i is always locally unhappy, independent of the boundary configuration. To do this we

need $|\Delta M_i^{local}(\sigma_i, \sigma_{\partial i})| > 2h(m - 1/N)$ that could happen with a nonzero probability as long as $2h(m - 1/N) < 1$. Now consider configuration $(+1, -1, +1)$ where player i is always locally happy. Again we can avoid this solution by choosing $|\Delta M_i^{local}(\sigma_i, \sigma_{\partial i})| < 2h(m + 1/N)$, which is possible if $2h(m + 1/N) > 0$.

Then consider three configurations $(-1, +1, -1)$, $(-1, -1, +1)$, $(-1, -1, -1)$ where player j could be locally happy or unhappy. In each case we can eliminate the solution by choosing $|\Delta M_j^{local}(\sigma_j, \sigma_{\partial j})| < 2h(m + 1/N)$ or $|\Delta M_j^{local}(\sigma_j, \sigma_{\partial j})| > -2h(m + 1/N)$. This can be done with a nonzero probability if $2h(m + 1/N) > 0$.

There remain two configurations $(+1, -1, -1)$ and $(+1, +1, +1)$ where player k could be locally happy or unhappy. Again we can choose $|\Delta M_k^{local}(\sigma_k, \sigma_{\partial k})| < 2h(m + 1/N)$ or $|\Delta M_k^{local}(\sigma_k, \sigma_{\partial k})| > -2h(m + 1/N)$ to eliminate $(+1, -1, -1)$ if $2h(m + 1/N) > 0$. And we choose $|\Delta M_k^{local}(\sigma_k, \sigma_{\partial k})| > 2h(m - 1/N)$ to eliminate $(+1, +1, +1)$ if $2h(m + 1/N) < 1$. Therefore, we have a nonzero probability $P_{nosolution}(\Omega)$ which in the thermodynamics limit gives an exponentially small probability of having solution $P_{solution} < (1 - P_{nosolution}(\Omega))^{N_\Omega}$.

-
- [1] L. E. Blume, W. A. Brock, S. N. Durlauf, and Y. M. Ioannides, in *Handbook of Social Economics*, edited by J. Benhabib, A. Bisin, and M. Jackson, eds., Amsterdam: North Holland, 2010 (forthcoming).
 - [2] S. Durlauf, and Y. Ioannides, *Annual Review of Economics* **2**, 451-478, 2010.
 - [3] O. Özgür, in *Handbook of Social Economics*, edited by J. Benhabib, A. Bisin, and M. Jackson, eds., Amsterdam: North Holland, 2010 (forthcoming).
 - [4] J. von Neumann, and O. Morgenstern, *The Theory of Games and Economic Behavior*. Princeton University Press, 1944.
 - [5] J. F. Nash, *Proc. Natl. Acad. Sci.* **36**, 48-49, 1950.
 - [6] J. F. Nash, *The Annals of Mathematics* **54**, 286-295, 1951.
 - [7] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, Cambridge University Press, 2007
 - [8] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, *Symp. on Theory of Computing*, pp 71-78, 2006.
 - [9] C. Daskalakis, and C. H. Papadimitriou, *Electronic Colloquium in Computational Complexity*,

- 139, 2005.
- [10] X. Chen, and X. Deng, Electronic Colloquium on Computational Complexity, 134, 2005.
 - [11] X. Chen, and X. Deng, Proceedings of FOCS, 261-272, 2006.
 - [12] C. H. Papadimitriou, J. Comp. Syst. Sci. 48, 498-532, 1994.
 - [13] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, Communications of the ACM **52**, pp 89-97, 2009.
 - [14] C. Daskalakis, Computer Science Review **3**, 87-100, 2009.
 - [15] I. Gilboa, and E. Zemel, Econ. Behav. **1**, 80-93, 1989.
 - [16] V. Conitzer, and T. Sandholm, Games and Economic Behavior **63**, 621-641, 2008.
 - [17] G. Gottlob, G. Greco, and F. Scarcello, Journal of Artificial Intelligence Research **24**, 357-406, 2005.
 - [18] A. X. Jiang, and M. Safari, In Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1 (AAMAS '10), Vol. 1. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 199-206, 2010.
 - [19] M. Mézard and A. Montanari, *Information, Physics, and Computation*, Oxford University Press, Oxford, 2009.
 - [20] M. Mézard, and G. Parisi, Eur. Phys. J. B **20**, 217, 2001.
 - [21] M. Mézard, G. Parisi, and R. Zecchina, Science **297**, 812 (2002).
 - [22] M. Mézard, and G. Parisi, J. Stat. Phys 111 (2003) 1.
 - [23] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian and L. Zdeborová, Proc. Natl. Acad. Sci. **104**, 10318 (2007).
 - [24] U. Horst, and J. Scheinkman, Journal of Economic Theory, **130** 44-77, 2006.
 - [25] U. Horst, and J. Scheinkman, To appear in Journal of Mathematical Economics.
 - [26] A. Bisin, U. Horst, and O. Özgür, J. Econ. Theory **127**, 74-116, 2006.
 - [27] Y. Bramoullé, and R. Kranton, Journal of Economic Theory **135**, 478-494 (2007).
 - [28] A. Galeotti, S. Goyal, M. Jackson, F. Vega-Redondo, and L. Yariv, The Review of Economic Studies (2009).
 - [29] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, IEEE Trans. Inf. Theory **47**, 498 (2001).
 - [30] J. Chavas, C. Furtlehner, M. Mezard, and R. Zecchina, J. Stat. Mech. P11016, 2005.
 - [31] A. Braunstein, and R. Zecchina, Phys. Rev. Lett. **96**, 030201, 2006.
 - [32] M. A. Nowak and R. M. May, Nature **359**, 826, 1992.

- [33] P. La Mura, Game networks, in Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI), pages 335-342, 2000.
- [34] M. Kearns, M. Littman, and S. Singh, Proceedings of the Conference on Uncertainty in Artificial Intelligence, 253-260, 2001.
- [35] D. Koller, and B. Milch, Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01), p 1027-1034, 2001.
- [36] L. E. Ortiz, and M. Kearns, Proceeding of NISP 2002.
- [37] D. Vickrey, and D. Koller, Proceedings of the National Conference on Artificial Intelligence (AAAI), 2002.
- [38] C. Daskalakis, and C. H. Papadimitriou, Proceedings of the 7th ACM conference on Electronic commerce, p.91-99, June 11-15, 2006.
- [39] J. Berg and A. Engel, *Phy. Rev. Lett.* **81**, 4999, 1998.
- [40] G. Szabo and G. Fath, *Phys. Rep.* **446**, 97, 2007.
- [41] C. Daskalakis, A. G. Dimakis, and E. Mossel, to appear in *Annals of Applied Probability*; arXiv:math/0703902v3 [math.PR].
- [42] L. Dall'Asta, P. Pin, and A. Ramezanzpour, *Phys. Rev. E* **80**, 061136, 2009.
- [43] J. -P. Onnela, and F. Reed-Tsochas, *Proc. Natl. Acad. Sci.* **43**, 18375-18380, 2009.
- [44] C. Di, A. Montanari, and R. Urbanke, *International Symposium on Information Theory*, IEEE, 2004.
- [45] T. Mora, and M. Mézard, *J. Stat. Mech.* 10:P10007, 2006.
- [46] P. Sulc, and L. Zdeborová, *J. Phys. A: Math. Theor.* **43**, 285003, 2010.